

INSTRUCTIONS FOR ORDERING EXECUTION IN PIPELINED PROCESSES

BACKGROUND OF THE INVENTION

Technical Field

5 This invention relates to pipelining processes in a multiprocessor computing environment. More specifically, the invention relates to a method and system for improving throughput based upon ordering constraints for shared memory operations.

Description Of The Prior Art

10 Multiprocessor systems contain multiple processors (also referred to herein as CPUs) that can execute multiple processes or multiple threads within a single process simultaneously in a manner known as parallel computing. In general, multiprocessor systems execute multiple processes or threads faster than conventional single processor systems, such as personal computer, that execute programs sequentially. The actual performance advantage is a function of a number of factors, including the degree to which parts of a multithreaded process and/or
15 multiple distinct processes can be executed in parallel and the architecture of the particular multiprocessor system. The degree to which processes can be executed in parallel depends, in part, on the extent to which they compete for exclusive access to shared memory resources.

20 Shared memory multiprocessor systems offer a common physical memory address space that all processors can access. Multiple processes therein, or multiple threads within a process, can communicate through shared variables in memory which allow the processes to read or write to the same memory location in the computer system. In order to increase operating efficiency in a multiprocessor system it is important to increase the speed by which a processor executes a program. One way to achieve this goal is to execute more than one operation at the same time. This approach is generally referred to as parallelism. A known technique for supporting parallel
25 programming and to manage memory access operations in a multiprocessor is pipelining. Pipelining is a technique in which the execution of an operation is partitioned into a series of

independent, sequential steps called pipeline segments. Each segment in the pipeline completes a part of an instructions, and different segments of different instructions may operate in parallel. Accordingly, pipelining is a form of instruction level parallelism that allows more than one operation to be processed in a pipeline at a given point in time.

5 In a cache-coherent system, multiple processors see a consistent view of memory. Several memory-consistency models may be implemented. The most straightforward model is called sequential consistency. Sequential consistency requires that the result of any execution be the same as if the accesses executed by each processor were kept in order and the accesses among different processors were interleaved. The simplest way to implement sequential

10 consistency is to require a processor to delay the completion of any memory access. However, sequential consistency is generally inefficient. Figs. 1 a-c outline the process of adding a new element 30 to a data structure 5 in a sequential consistency model. Fig. 1a is an illustration of a sequential consistency memory model for a data structure prior to adding or initializing a new element 30 to the data structure 5. The data structure 5 includes a first element 10 and a second element 20. Both the first and second elements 10 and 20, respectively, have three fields 12, 14 and 15, and 22, 24 and 26. In order to add a new element 30 to the data structure 5 such that the CPUs in the multiprocessor environment could concurrently search the data structure, the new element 30 must first be initialized. This ensures that CPUs searching the linked data structure do not see fields in the new element filled with corrupted data. Following initialization of the

15 new element's 30 fields 32, 34 and 36, the new element may be added to the data structure 5. Fig. 1b is an illustration of the new element 30 following initialization of each of it's fields 32, 34 and 36, and prior to adding the new element 30 to the data structure 5. Finally, Fig. 1c illustrates the addition of the third element to the data structure following the initialization of the fields 32, 34 and 36. Accordingly, in a sequential consistency memory model execution of each

20

25 step in the process must occur in a pre-specified order.

The process of Figs. 1a-c is only effective on CPUs that use a sequentially consistent memory model. For example, the sequential memory model may fail in weaker memory models

where other CPUs may see write operations from a given CPU happening in different orders.

Fig. 2 is an illustration of a weak memory-consistency model for adding a new element to a data structure. In this example, the write operation to the new element's 30 first field 32 passes the write operation to the second element's 20 next field 22. A CPU searching the data structure may see the first field 32 of the third element 30, resulting in corrupted data. The searching CPU may then attempt to use the data ascertained from the field 32 as a pointer, and most likely this would result in a program failure or a system crash. Accordingly, it is desirable to place some form of a memory barrier instruction to be executed prior to storing a pointer from the second element in the data structure to the new element in the data structure.

Fig. 3 is a block diagram 40 illustrating the segregation of instructions into groups, wherein one group of instructions occurs before the memory barrier and another group of instructions occurs after the memory barrier. This diagram follows the linked data structure example of Figs. 1 and 2. There are essentially four levels of operation. The first level includes the following operations: storing a NULL pointer into the new element's first field 42, storing the character string "IJKL" into the new element's second field 44, and storing the number "9012" into the new elements third field 46. Following this group of write operations, a memory barrier 50 is executed. The memory barrier ensures that each of the write operations 42, 44 and 46 occur prior to any other computations. Following the execution of the memory barrier 50 and the execution of the write operations 42, 44 and 46, the address of the second element may be computed 52. Step 52 is a local memory operation, and it may involve a plurality of write operations to the CPU's local memory. Finally, following step 52, a pointer to the new element is stored in the second element's first field 54. Although the memory barrier instruction 50 prevents the memory write operations 42, 44 and 46 from appearing to have occurred later than memory write operation at 54, it needlessly prevents the write operations in 42, 44 and 46 from appearing to have occurred later than computation of address 52. Accordingly, the prior uses of memory barrier instructions as shown in Fig. 3 results in an inefficient use of the CPU's resources resulting in a delayed execution of the program.

Fig. 4 is a block diagram 60 similar to the example shown in Fig. 3 without the memory barrier instruction. This diagram follows the linked data structure example of Figs. 1 and 2. In this example, the memory barrier instruction 50 is removed, and as such there are two levels of operation. The first level includes the following operations: storing a NULL pointer into the new element's first field 42, storing the character string "IJKL" into the new element's second field 44, and storing the number "9012" into the new elements third field 46. Following the write operations of 42, 44 and 46, the address of the second element may be computed 52. At the same time, a pointer to the new element is stored in the second element's first field 54. The removal of the memory barrier instruction allows the address of the second element to be computed 52 at the same time as storing a pointer to the new element in the second element's first field 54. The removal of the memory barrier instruction increases the efficiency of operation of the program. However, there may temporarily be corrupted data in the new element. Accordingly, there is a need for an efficient pipelining model that maintains data integrity while improving operating efficiency.

One programming model that allows a more efficient implementation is synchronization. A program is synchronized if all access to shared data is ordered by synchronized operations. In addition to synchronizing programs, there is also a need to define the ordering of memory operations. There are two types of restrictions on memory orders, write barriers and read barriers. In general, barriers act as boundaries, forcing the processor to order read operations and write operations with respect to the barrier. Barriers are fixed points in a computation that ensure that no read operation or write operation is moved across the barrier. For example, a write barrier executed by a processor A ensures that all write operations by A prior to the write barrier operation have completed, and no write operations that occur after the write barrier in A are initiated before the barrier operation. In sequential consistency, all read operations are read barriers, and all write operations are write barriers. This limits the ability of the hardware to optimize accesses, since order must be strictly maintained. The typical effect of a write barrier is to cause the program execution to stall until all outstanding writes have completed, including the delivery of any associated invalidations.

In an attempt to increase performance, it has become known to reorder execution of instructions. However, in reordering instructions special synchronizing instructions are required in order to specify to the CPU which accesses may not be reordered. Accordingly, there is a need for a computer system comprising multiple processors for maximizing CPU performance by placing constraints on shared memory access while removing constraints on non-shared memory accesses.

SUMMARY OF THE INVENTION

It is therefore an object of the invention to provide a method of maximizing CPU performance in a multiprocessor computer system. It is a further object of the invention to allow local memory operations to execute in an arbitrary order while providing constraints for shared memory operations.

A first aspect of the invention is a method of specifying ordering of execution of sets of instructions. Multiple sets of instructions are provided in which ordering of execution is maintained through implementation of registers, assignment of sequence numbers, or a hierarchical ordering system. First, second and third sets of instructions are provided. The third instruction set specifies the order of execution between the first and second sets of instructions. The third instruction set requires that the execution of the first instruction set reach a specified state of execution before the execution of the second instruction set reaches a specified state of execution. The specified state of execution of the first and second instruction sets is preferably, but not necessarily, selected from the group consist of: committing instruction execution, initiating memory access, completing a memory access, initiating an I/O access, completing an I/O access, and completing instruction execution.

A second aspect of the invention is a processor for use in a multiprocessor computer system which includes instructions for ordering operating constraints within a computer processing system. First, second, and third sets of instructions are provided. The third

instruction set is a manager to ensure the system maintains order of execution between the first and second sets of instructions, and that the first instruction set reach a specified state of execution before the second instruction set reaches a specified state of execution. The specified state of the first and second instruction is preferably, but not necessarily, selected from the group consisting of: committing instruction execution, initiating memory access, completing a memory access, initiating an I/O access, completing an I/O access, and completing instruction execution.

A third aspect of the invention is an article comprising a computer-readable signal bearing medium with multiple processors operating within the medium. The article includes a manager for scheduling shared memory operations. The manager utilizes multiple sets of instructions and specifies ordering of the instructions. The ordering of instructions is preferably, but not necessarily, selected from the group consisting of: utilizing a pair of CPU registers and assigning sets of instructions to each register, assigning sequence numbers to sets of instructions, and placing a range of instructions into a hierarchical ordering system.

Other features and advantages of this invention will become apparent from the following detailed description of the presently preferred embodiment of the invention, taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1a is a block diagram of a prior art data structure at an initial state.

FIG. 1b is a block diagram of a prior art data structure with a new element initialized.

FIG. 1c is a block diagram of a prior art data structure with a new element appended to a list.

FIG. 2 is a block diagram of a prior art data structure of a weak memory-consistency model.

FIG. 3 is a prior art block diagram illustrating use of a memory barrier for appending a data structure.

FIG. 4 is a prior art block diagram illustrating appending a data structure without a memory barrier.

5 FIG. 5 is a block diagram illustrating appending a data structure with implementation of special instructions according to the preferred embodiment of this invention, and is suggested for printing on the first page of the issued patent.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Overview

10 In a shared memory multiprocessor system it is essential that multiple processors see a consistent view of memory. CPUs that use a weak memory consistency model generally use memory barrier instructions to force the order of write operations. However, memory barrier instructions place constraints on the order of memory writes performed by all instructions. These constraints place artificial limits on the amount of performance increase that a CPU may attain
15 by reordering operations. In many algorithms, there are local memory operations that may be ordered arbitrarily wherein only certain global memory access operations need be carefully ordered. Accordingly, it is desirable and efficient to implement a method that allows selected write operations to global memory to be properly ordered, while allowing a CPU full freedom to reorder local write operations as needed to optimize CPU performance.

20 Technical Background

 In general, neither the CPU nor the compiler can distinguish between local and global memory operations. Nor can the CPU or the compiler determine which global memory operation access must be ordered. It is therefore necessary for the programmer to indicate the ordering by use of special compiler directives which would cause the compiler to insert special assembly
25 language instructions, or by inserting special assembly language instructions. Such instructions

explicitly indicate which write operations are to be placed in a specific order of operation and which write operations may occur arbitrarily.

Fig. 5 refers to a block diagram 90 of the process of ordering global memory operations. This process provides for a CPU to explicitly indicate which write operations are to be specifically ordered, while allowing the remainder of the write operations to occur at any time. The flow diagram illustrates the ordering of global memory operations. Similar to the block diagram of Figs. 3 and 4 and the data structure list example of Figs. 1 and 2, this model maintains an ordering of the processes prior to linking the new element 30 to the second element 20 of the data structure 5. As shown in Fig. 5, the process of writing to each of the field in the new element may occur at any time prior to storing a pointer from the second element 20 to the new element 30. Step 42 references the process of conducting a write operation to the first field of the new element, step 44 references the process of conducting a write operation of the second field of the new element, and step 46 references the process of conducting a write operation to the third field of the new element. Each of the write operations is monitored by a special instruction 82, 84 and 86, respectively. The independent special instruction associated with each write operation forces the write operation to precede step 54, which is the process of storing a pointer to the new element in the second element's first field 22. In conjunction with the write operation steps 42, 44, and 46, and the associated special instructions 82, 84, and 86, this preferred embodiment allows the local memory operations 52 to occur at any time prior to the process of storing a pointer to the new element in the second element's first field 54. As shown in Fig. 5, the local memory operation may occur in conjunction with the memory write operations 42, 44, and 46 or the special instruction operations 82, 84 and 86. Accordingly, the implementation of the special instructions in conjunction with the removal of the memory barrier that was shown in Fig. 4, explicitly indicates which write operations are to be conducted in a specified order and which write operations may be conducted in an arbitrary order.

Pseudocode for the special instructions that explicitly indicates which write operations are to be executed in a specific order is as follows:

1. Store a NULL pointer into the new element's first field.
2. Store the character string "IJKL" into the new element's second field.
3. Store the number 9012 into the new element's third field.
4. Execute a special instruction that forces the write in step 1 to precede that in step 8.
5. Execute a special instruction that forces the write in step 2 to precede that in step 8.
6. Execute a special instruction that forces the write in step 2 to precede that in step 8.
7. Compute the address of the second element (which could involve many write operations to local memory).
8. Store a pointer to the new element into the second element's first field.

The pseudocode outlined above demonstrates a process that allows any local memory write operations to proceed at any time, *i.e.* either before, during or after the process of storing data in the fields of the data structure, but prior to the process of storing a pointer from the existing data structure to the new data structure element. Accordingly, the flexibility of allowing local memory write operation to proceed at any time prior to the process of establishing a pointer to the new element of the data structure provides the CPU the freedom to optimize use of its internal resources.

Furthermore, as is shown in Fig. 5, there are essentially three levels of operation that occur in linking the new element 30 to the second element 20. The first level 100 is the write operations to each of the respective fields of the new element 30. The second level 200 is the special instruction associated with each of the write operations. The third level 300 is the process of storing a pointer to the new element 30 into the second element's first field 22. The local memory operations, as referenced by 52, may occur in conjunction with the processes in the first level 100 or the second level 200. Accordingly, the implementation of the special

instructions in conjunction with removal of the memory barrier illustrated in Fig. 3 actually speeds up the process of adding the new element while maintaining the integrity of the data.

There are several methods of designing instructions to order global memory operations as shown in Fig. 5. One method is to set aside a pair of registers of a CPU for assigning address instructions. The register is resident in a CPU and an external CPU cannot access the registers of another CPU. The first register contains a first instruction address. The second register contains a second instruction address. The execution of a special third instruction will reference the instructions indicated by the first and second registers. The third instruction specifies ordering between the first and second instructions. The remainder of the instructions that are not referenced in the registers can occur at any time during the execution of the write operations. In general, the global memory operations are the instructions that are referenced in the registers. The third instruction ensures that the first instruction's execution reach a predefined state prior to the execution of the second instruction reaching a predefined state. Examples of the predefined state are: committing instruction execution, initiating a memory access, completing a memory access, initiating an I/O access, completing an I/O access, and completing instruction execution. Both the first and second instruction may have the same predefined state, or the states may be separately defined for the different instructions or for different groups of the instructions. Accordingly, the assignment of instruction addresses into a pair of registers is just one embodiment of how the process of the preferred embodiment may be implemented.

Another method of implementing the process of the preferred embodiment is to control the order in which memory write operations are flushed from the write buffer. Each write buffer entry is assigned a sequence number for identifying the sensitivity of the entry. The sequence number indicates the order in which entries must be flushed to memory. Alternatively, the sequence number may indicate that the corresponding entry may be flushed to memory at any time and is not order dependent. The hardware that flushes the write buffer would then have the information required to flush the buffers in the hierarchical ordering provided. As in the case of registers, an external CPU cannot see the write buffer. The following table illustrates the

appearance of a CPU's write buffer in conjunction with the eight step pseudocode of the preferred embodiment:

	Address	Data	Sequence
5	First field of new element	NULL pointer	1
	Second field of new element	"IJKL"	1
	Third field of new element	9012	1
	Local variables used to compute address of second element	Local stack address	Don't Care
10	Another local variable	Local stack address	Don't Care
	First field of second element	Pointer to new element	2

The "don't care" values are associated with local memory operations, and allows the CPU to optimize performance. Alternatively, the "don't care" values may be indicative that the associated instruction is not sensitive to order of execution, and that conducting the associated operation out of order will not affect the integrity of the associated data. The filling out of the data structure can occur at any time prior to the establishment of a pointer to the new element of the data structure. The remainder of the processes are conducted in numerical order governed by the associated sequence number.

In a further embodiment of the invention, the sequence numbers may be provided statically, wherein the sequence numbers are encoded directly into the instruction. Alternatively, the sequence numbers may be dynamically encoded. An example of dynamically encoding sequence number is to read the instruction sequence numbers out of the CPU registers. Accordingly, the alternative method of providing sequence numbers for the associated instructions may provide a more efficient performance since the order of operation of the instructions may be adapted to different circumstances.

Another example of implementing instructions indicating specificity of order of global memory operations is to place a range of instructions into a hierarchical ordering system. In this

format, a range of instructions are placed into a group, and multiple other instructions or a range of instructions may be placed into other groups. A special instruction is implemented to ensure that the hierarchical ordering of the groups is maintained. This ensures that the integrity of the data written to the data structure is not temporarily corrupted.

- 5 The pseudocode for the placement of a range of instructions into a hierarchical ordering system is as follows:
1. Store a NULL pointer into the new element's first field.
 2. Store the character string "IJKL" into the new element's second field.
 3. Store the number 9012 into the new element's third field.
 - 10 4. Execute a special instruction that groups the write operations in steps 1, 2 and 3.
 5. Execute a special instruction that forces the write operations in the group indicated by Step 4 to precede that in Step 7.
 6. Compute the address of the second element (which could involve many write operations to local memory).
 - 15 7. Store a pointer to the new element into the second element's first field.

The hierarchical ordering system is advantageous where elements with many field are being inserted into a list while allowing concurrent readers. Accordingly, the hierarchical ordering system provides for a special instruction to place multiple write operations into a grouping, wherein another instruction indicates order of operation of assigned groupings.

20

Alternative Embodiments

It will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without departing from the spirit and scope of the invention. In particular, a second embodiment is a microprocessor for use in a multiprocessor computer system. The microprocessor contains

registers and instructions as described above, implemented using well known skills in creating instruction sets and register assignments for processors. More specifically, a first instruction is provided to allow local memory operations to occur in an arbitrary order, and a second instruction is provided to place constraints on shared memory operations. The first instruction is indicative of the absence of an instruction for local memory operations. A third instruction is provided to manage the order of execution of the first and second instructions. Execution of the second instruction is responsive to the first instruction reaching a specified state of execution. Examples of the state of execution are: committing instruction execution, initiating an I/O access, completing an I/O access, and completing an instruction execution. There are several alternative components used in implementing the hierarchical ordering of instructions, including storing the first and second instructions on separate registers resident in a CPU, assigning sequence numbers to the instructions for specifying the order of execution, and implementing a manager to place a range of instructions in a hierarchical order. Accordingly, the scope of protection of this invention is limited only by the following claims and their equivalents.